

Naval Research Laboratory

Washington, DC 20375-5320



NRL/MR/5620--94-7612

AD-A285 406



System Design, Implementation, and Evaluation of the Optical Broadband Correlator

111-1994
IC
D

R. D. GRIFFIN

*Advanced Concepts Branch
Optical Sciences Division*

September 20, 1994

Approved for public release; distribution unlimited.

2018
94-31906



REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project 0704-0188, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)			2. REPORT DATE		3. REPORT TYPE AND DATES COVERED		
			September 20, 1994				
4. TITLE AND SUBTITLE				5. FUNDING NUMBERS			
System Design, Implementation, and Evaluation of the Optical Broadband Correlator				PE - 62314N			
6. AUTHOR(S)							
R. D. Griffin							
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER			
Naval Research Laboratory Washington, DC 20375-5320				NRL/MR/5620-94-7612			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSORING/MONITORING AGENCY REPORT NUMBER			
Office of Naval Research Ballston Tower One, 800 North Quincy Street Arlington, VA 22216-5660							
11. SUPPLEMENTARY NOTES							
12a. DISTRIBUTION/AVAILABILITY STATEMENT				12b. DISTRIBUTION CODE			
Approved for public release; distribution unlimited.							
13. ABSTRACT (Maximum 200 words)							
This report describes the design, implementation, and evaluation of an acousto-optical wide band correlator that is integrated into a digital signal processing system testbed. The report compares the correlator's measured performance with various theoretical performance measures, one of which compares the optical system with an equivalent efficient digital correlator. This measured performance of the optical system was 20 to 70 times that of a VAX 6410 using an FFT correlation algorithm and a vector processor, even though the electronic interface system limited the performance of the optical system to less than one-half percent of its potential.							
14. SUBJECT TERMS				15. NUMBER OF PAGES			
Acousto-optics Correlator Matched filter				29			
16. PRICE CODE							
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE		19. SECURITY CLASSIFICATION OF ABSTRACT		20. LIMITATION OF ABSTRACT	
UNCLASSIFIED		UNCLASSIFIED		UNCLASSIFIED		UL	

— CONTENTS

INTRODUCTION.....	1
ACOUSTO-OPTICAL SIGNAL PROCESSING.....	2
Band Matching.....	3
Performance Measures.....	4
Input and Output Bandwidth Requirements.....	10
DESIGN GOALS.....	11
SYSTEM ARCHITECTURE.....	13
Hardware.....	13
Optical System.....	14
Electronic System.....	16
Software.....	18
RESULTS.....	19
I/O Bandwidth Performance.....	19
Computational Performance Measurements.....	20
Dynamic Range Measurements.....	22
CONCLUSION.....	24
ACKNOWLEDGMENTS.....	25
REFERENCES	25

Accesion For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and / or Special

A-1

SYSTEM DESIGN, IMPLEMENTATION, AND EVALUATION OF THE OPTICAL BROADBAND CORRELATOR

INTRODUCTION

Acousto-optical (AO) techniques and hardware have been described extensively in the literature,^{1,2} but relatively little has been reported on the integration of AO technology into actual signal processing systems.³ We have developed a space-integrating acousto-optical correlator system that is integrated as a matched filter processor (or correlation receiver) into a testbed digital signal processing system. In this report we describe the design goals of the system, describe some of the issues associated with utilizing high-performance, specialized processors; derive several performance measures; and report the results of performance measurement tests.

Acousto-optical technology has the potential to provide a 100- to 1000-fold improvement in matched-filter processing power over current digital systems, with potentially lower volume and power requirements. The optical component of our system has the potential to provide the equivalent of 8-10 GFLOPS of processing power, although with its current electronic interface, less than one-half percent of this potential is harnessed. Nevertheless, it performed 20 to 70 times faster than a VAX 6410 that used a vector processor and an optimized FFT correlation routine. Thus, it has the potential to run approximately ten-thousand times faster than the VAX. The optical system and electronic interface occupy 1.5 ft³ and consume approximately 200 W of power. This yields a potential system figure of merit of 27-33 GFLOPS/ft³-kW. Two-dimensional optical techniques have the potential to enhance the computation rate by a factor of 25 to 100, and improved electronics and packaging can further improve the figure of merit. Integrating such a powerful computational engine into a digital signal processing system is a significant challenge due to the large input and output data bandwidths, and to the necessary conversions between the analog and digital domains.

The matched filter concept⁴ is fundamental to a wide variety of decision making systems, whether for communicating information, for detecting signals or objects, or for estimating some parameter in a radar, sonar, pattern recognition, or signal classification system. The matched filter is a linear filter with an impulse response that has the same form as the signal of interest except for a reversal in time; it is equivalent to a correlation receiver. Typically, the output of a matched filter is compared with a reference level, and the presence of a signal (a detection) is declared whenever the output level exceeds the reference level. The location in time of the detection, the matched filter output level, and other features of the matched filter output are often used to estimate various parameters of interest such as target range, strength, and speed. The correlation operation can also be used as a building block for more complex operations.⁵ An example of recent interest is wavelet transform processing,⁶ wherein a signal is crosscorrelated with time-dilated and time-shifted versions of a wavelet function.

Implementing a matched filter with a correlation receiver is useful because one correlator operating in *compressed time* can replace a large bank of matched filters operating in real time: if the correlator takes T/ϵ seconds to process T seconds of signal, then ϵ matched filters can be replaced by one correlator. The disadvantage with this approach is that the signal waveform must be stored so that it can be processed

ε times. This is not a problem with digital systems, since the waveform has already been sampled, digitized, and stored. The primary limitation with current digital correlation receivers is the amount of time compression ε they can provide for a given budget that may include cost, volume, cooling, and power constraints.

Acousto-optical technology is a natural fit¹ to the convolution or correlation problem, but harnessing the power of an acousto-optical correlation receiver requires some effort, since there are several constraints or difficulties that limit the application of the technology to the matched filter problem. One of the most severe constraints is the average bandwidth required to keep an acousto-optical correlator supplied with data, and the average bandwidth required to receive its output—bandwidths that can easily reach hundreds or even thousands of megabytes per second. Rather than build a system with full bandwidth capability, we chose an application where the input bandwidth is small and the output bandwidth can be reduced.

Since AO modulator technology is analog and high bandwidth, a second difficulty is obtaining digital-to-analog converters (DAC's) and analog-to-digital converters (ADC's) with enough resolution at the requisite speeds. Until recently, this constraint was enough to limit AO technology to a handful of applications where analog signal processing could be used or where only a few bits of resolution were required. With advances in digital technology, these constraints are rapidly disappearing, although they are still a significant issue in the design, construction, and testing of systems like the one described here.

A third concern is the accuracy or fidelity of the operations provided by the acousto-optical system. The accuracy depends on parameters such as the linearity, dynamic range, resolution, and signal-to-noise ratio of the various optical and electronic devices in the system.

A fourth concern is that the acousto-optical nature of the matched filter system be transparent to the user. That is, a user should be able to access the AO correlator in exactly the same manner as a digital correlator, and the results he obtains from either system should ideally be identical. It is a relatively simple matter to provide the first feature, but the second feature is not possible: rather, the degree to which the results will be identical must be set as a goal and incorporated into the overall design.

ACOUSTO-OPTICAL SIGNAL PROCESSING

Before we describe our design goals and the system architecture, we first discuss some signal processing aspects of correlation or convolution with acousto-optical modulators.¹ The basic optical layout for a space-integrating correlator is shown in Fig. 1, where two acousto-optical modulators or cells are shown illuminated by a collimated beam of coherent light. The width of the beam is called the optical aperture. Two counter-propagating acoustical waves are generated in the acousto-optical cells by rf signals labeled Replica and Signal. Each traveling wave diffracts a portion of the incident light. The second AO cell is adjusted such that it mainly diffracts the diffracted light from the first cell. The correlation signal results from optical heterodyne detection of the doubly-diffracted light mixed with the undiffracted light in the photodiode detector. The details of the optical hardware are described in a later section. We chose a space-integrating architecture because the replica waveforms in our intended application fit within a standard Bragg cell, and because it was more useful to produce the output directly in the time domain. A time-integrating correlator would require a photodetector array, which has a lower dynamic range because of its non uniform response across the array. If the replica waveforms are longer than the Bragg cell, the time-integrating architecture is a necessity. Parts of the discussion that follows are unique to the space-integrating architecture, although similar concepts would apply to the time-integrating architecture.




Fig. 1.—Basic optical layout for correlation processing.

Band Matching

Acousto-optical cells typically operate at frequencies in the MHz to GHz range depending on the type of cell. If the spectrum of the input signal to an acousto-optical system does not lie in the band of the AO cell, the band of the input signals must be matched to that of the AO cell. Sometimes a signal, e.g., in radar, will have a band that naturally matches the band of a particular cell, or there is enough freedom in the system requirements to design the band of the signal to match that of a suitable AO cell. However, many signals of interest have bandwidths and center frequencies substantially different from those of available AO cells. If the bandwidth is a close match, but the center frequency is not, the signal may be band-shifted by mixing with a suitable carrier followed by filtering. If the signal bandwidth is substantially less than that of the AO cell, the signal must be time-compressed to increase the bandwidth. Time compression requires additional processing, but in the case of large compression factors, it can lead to substantial performance gains since multiple time-compressed signals may be processed during the duration of a single uncompressed signal. For example, correlation of two signals that are several seconds long but with bandwidths approximately one-millionth that of the AO cell can be completed in several microseconds.

In practical terms, time-compression requires that the signals of interest be digitized and stored in memory for later readout at a rate commensurate with the center frequency and bandwidth of the acousto-optical cells. The concept is illustrated in Fig. 2(a)-(c), which shows in (a) the bandwidth W and center frequency f_0 of the original signal sampled at a rate f_S . In (b) the original signal has been shifted by an amount $-\delta f$ to a lower center frequency f'_0 . If necessary, the signal is then resampled⁷ via an interpolation/decimation scheme by a factor (L/M) , where L and M are integers, to a new sample rate f'_S . In (c), the sample rate is changed to f_{CLK} which produces a time-compression factor ϵ given by

$$\epsilon = \frac{f_{CLK}}{f_S} = \left(\frac{M}{L} \right) \frac{f_{CLK}}{f'_S}. \quad (1)$$

Thus, with this value of ε , and δf chosen to satisfy

$$f_{cell} = \varepsilon f'_0 = \varepsilon(f_0 - \delta f), \quad (2)$$

the time-compressed signal will have bandwidth W_{cell} and center frequency f_{cell} . Note that if f_{CLK} is not fixed and has sufficient variability, the resampling operation can be eliminated by choosing the compression factor according to

$$\varepsilon = \frac{W_{cell}}{W}. \quad (3)$$

A more general constraint is that the compression factors be chosen so that the band of the compressed signal lies within the band of the AO cells. In practice, this means that L and M may be chosen so that the center frequency of the compressed signal is *near* f_{cell} , and that the compressed bandwidth be no greater than W_{cell} .

Performance Measures

To have some measure of the performance of the acousto-optical correlator we define several computation rates, and for comparison purposes we compute the computational performance required of an equivalent digital processor. To do this we must first analyze the timing for a basic correlation sweep when implemented with a space-integrating acousto-optical correlator.

The diagram of Fig. 3 shows the relative positions of the reference and return waveforms within the optical aperture at the beginning of a correlation sweep. It also shows a second set of reference and return waveforms that will generate the next correlation function. The reference waveform duration is indicated by T_{REF} , the return or received waveform duration by T_{RCV} , and the optical aperture length is indicated by T_{OA} . The reference (or replica) waveform propagates to the right and the return waveform propagates to the left. The correlation sweep begins when the left-hand or trailing edge of the replica meets the left-hand or leading edge of the return, and it ends when the leading edge of the replica meets the trailing edge of the return. In the example shown in Fig. 3, the correlation function is generated for a range of T_{COR} seconds of delay or lag given by

$$T_{COR} = (T_{RCV} - T_{REF}), \quad (4)$$

where \bar{T}_{COR} is the actual duration of the correlation function in real time, and accounts for the effect of the counter-propagating waveforms. One can also see from the figure that the longest return waveform T_{RCV}^{max} that can be processed during one correlation sweep is given in compressed time by

$$T_{RCV}^{max} = 2T_{OA} - T_{REF}. \quad (5)$$

Using a return waveform shorter than this is inefficient since part of the optical aperture is not used during the correlation; as such, we will frequently make use of this relation in what follows without explicit mention. Note that the reference signal must fit within the optical aperture, which implies that its duration T'_{REF} must satisfy

$$T'_{REF} \leq \varepsilon T_{OA}. \quad (6)$$




Fig. 2.—Time compression of a waveform. (a) Band of original waveform sampled at a rate f_s . (b) Shifted by an amount $-\delta f$. (c) Time compression effected by change of sample rate to f_{CLK} .

One can also see from Fig. 3 that the actual duration of the correlation function is just half the duration of the range swept out. This range compression effect is due to the counter propagation of the waveforms; it effectively doubles the bandwidth of the correlation function, which has implications for the analog-to-digital converters and other electronic circuitry used in the post processor.

A return waveform that is longer than T_{RCV}^{max} can be processed in sections, or it can be processed as a single waveform by correlating it against suitably delayed copies of the replica waveform. To satisfy system requirements we chose the former method, even though the latter method is more efficient. Our method of sectioning is equivalent to the *overlap-save* method⁸ for discrete Fourier transform convolution and correlation. Because of the overhead time required to execute software, set up counters, address registers, clocks, and so on, the system requires a delay between correlation sections. The effect of this

overhead time T_{OVR} is shown in Fig. 3 for the second set of reference and return waveforms. Thus, the correlation repetition time or period is given by

$$T_{RPT} = T_{RCV} + T_{OVR}. \quad (7)$$




Fig. 3.—Reference and return waveform positions relative to the optical aperture at the beginning of a correlation sweep. Also shows the effect of overhead on the positions of a second set of waveforms that will produce the next correlation sweep.

We now define several measures of computation performance. One of the simplest measures we call the raw computation rate. This is just the rate at which equivalent digital multiplications are done to produce one sample value of the correlation function. Since the reference and return waveforms are counter propagating, the correlation sample rate must be twice the input waveform sample rate, and since the number of multiplications done to produce one correlation function sample is the number of replica samples N_{REF} , the raw rate is given by

$$f_{raw} = \frac{N_{REF}}{1/(2f_{CLK})} = 2f_{CLK}N_{REF}. \quad (8)$$

This raw rate is the effective number of complex multiplications that would be done by a digital computer to produce one correlation function sample value using the straightforward definition of the discrete correlation function.

Another performance measure is the average number of correlation function samples produced per second. This is given by the ratio of the number of samples in a correlation function to the correlation period, that is,

$$f_{avg} = \frac{(2f_{CLK})(\frac{1}{2}T_{COR})}{T_{RPT}} = f_{CLK} \frac{T_{RCV} - T_{REF}}{T_{RCV} + T_{OVR}}. \quad (9)$$

The maximum value for the average rate occurs when we use T_{RCV}^{\max} , and set T_{OVR} to zero. Using a small value for T_{REF} defeats the purpose of the correlator, so that is not an option. Thus, we define the maximum average output rate as

$$f_{avg}^{\max} = 2f_{CLK} \frac{T_{OA} - T_{REF}}{2T_{OA} - T_{REF}} = 2f_{CLK} \frac{1 - \tau_{REF}}{2 - \tau_{REF}}, \quad (10)$$

where

$$\tau_{REF} = \frac{T_{REF}}{T_{OA}} \quad (11)$$

is the duration of the reference waveform relative to the optical aperture. Note that the maximum average rate is independent of the length of the optical aperture, and depends only on the clock rate and τ_{REF} . From Eq. (6) we see that τ_{REF} can only vary between 0 and 1.

We can now define an efficiency that tells us for a given value of T_{REF} what the effect of T_{OVR} is on the average computation rate. This efficiency η is just the ratio of the actual average rate to the theoretical maximum average rate, and is given by

$$\eta = \frac{f_{avg}}{f_{avg}^{\max}}. \quad (12)$$

The final performance measure compares the optical correlator with an equivalent digital processor. To do that we compute the number of floating point operations required of an efficient digital correlator to process a long return signal of duration T_{sig} , and then divide that number by the time required for the optical correlator to process the same signal. This definition allows each method to use its most efficient mode of processing. For our comparison the digital correlator includes three efficiencies. The first efficiency is introduced by using minimum-sampled⁹ baseband signals, which minimizes the number of samples that must be processed. A second efficiency is introduced by using the Fast Fourier Transform algorithm⁸ for sectioned correlations. The third efficiency comes from using an optimum FFT size¹⁰ based on the size of the replica waveform.

If we ignore the operations required to Fourier transform the replica waveform (since it need be done only once), the number of operations (complex multiplication-additions) required for the FFT correlation⁸ of one section of return waveform is

$$n_{op} = 2N_{opt} \left(\log_2 N_{opt} + \frac{1}{2} \right), \quad (13)$$

where N_{opt} is the optimum FFT size for the replica. This optimum size minimizes the number of computations required for the FFT assuming a return waveform that is much longer than the replica waveform. It is determined empirically and is a function of N_{REF} as shown in Table 1.

With the effects of overlap included, the number of sections required to correlate the entire return waveform is

$$n_{sect} = \frac{f_S T_{sig}}{N_{opt} - N_{REF}}. \quad (14)$$

where the term N_{REF} accounts for the samples in the overlap section. The total number of operations required to process the return waveform is $n_{op} n_{sect}$.

Table 1—Optimum Values of FFT Size for Correlation or Convolution

N_{REF}	N_{opt}	$\log_2 N_{opt}$
< 11	32	5
11 - 17	64	6
18 - 29	128	7
30 - 52	256	8
53 - 94	512	9
95 - 171	1024	10
172 - 310	2048	11
310 - 575	4096	12
575 - 1050	8192	13
1050 - 2000	16384	14
2000 - 3800	32768	15
3800 - 7400	65536	16
> 7400	131072	17

The time required for the optical correlator to process the same waveform is $n_{sect}^{AO} T_{RPT}$, where the number of return waveform sections required by the AO correlator is

$$n_{sect}^{AO} = \frac{T_{sig} / \varepsilon}{T_{RCV}^{max} - T_{REF}} = \frac{T_{sig}}{2\varepsilon(1 - \tau_{REF})T_{OA}}. \quad (15)$$

after substituting from Eq. (5). Thus, the required equivalent digital processing power in floating point operations per second is given by

$$f_{FLOPS} = \frac{n_{op} n_{sect}}{n_{sect}^{AO} T_{RPT}} = 4\varepsilon f_S \frac{(1 - \tau_{REF})}{(2 - \tau_{REF} + \tau_{OVR})} \frac{N_{opt} (\log_2 N_{opt} + \frac{1}{2})}{(N_{opt} - N_{REF})}. \quad (16)$$

where

$$\tau_{OVR} = \frac{T_{OVR}}{T_{OA}}. \quad (17)$$

This definition of equivalent processing power recognizes that each method may have different optimum parameters (such as the section length). It does not include the processing required to baseband the data, although this should be included if it is done for the digital processor but not for the AO processor.

Eq. (16) is only an estimate: a good engineering rule of thumb is to increase this number by a factor of four or five to account for the overhead in a typical digital machine.

A useful approximation to Eq. (16) can be made in the following way. Since the condition $T_{REF} < T_{OA}$ must always hold, the first ratio in the equation is never greater than one-half. An examination of Table 1 shows that for the larger values of N_{REF} , N_{opt} is a rough approximation to $N_{opt} - N_{REF}$; thus, we may replace the second ratio by $\log_2 N_{opt}$. Under these assumptions the equivalent processor power becomes

$$f_{FLOPS} \approx 2\epsilon f_S \log_2 N_{opt} = 2f_{CLK} \log_2 N_{opt}. \quad (18)$$

although this is something of an underestimate.

With the assumption of no overhead, Eq. (16) becomes

$$f_{FLOPS}^{effic} = 4f_{CLK} \frac{(1 - \tau_{REF})}{(2 - \tau_{REF})} \frac{N_{opt} \left(\log_2 N_{opt} + \frac{1}{2} \right)}{(N_{opt} - N_{REF})}. \quad (19)$$

Table 2—Computation Rates ($f_{CLK} = 80$ MHz, $N_{REF} = 1024$ and $N_{opt} = 8192$.)

τ_{REF}	f_{raw} (10^9 mult./s)	f_{avg}^{max} (MSPS)	f_{FLOPS}^{effic} (GFLOPS)
0	0	80.0	2.47
0.1	56.3	75.8	2.34
0.25	141	68.6	2.12
0.33	186	64.2	1.98
0.5	282	53.3	1.65
0.67	377	39.7	1.22
0.75	422	32.0	0.987
1.0	563	0	0

The various performance measures are compared in Table 2 for the value of f_{CLK} used in our system and for various replica durations. For the raw computation rate of Eq. (8), the multiplication effected is equivalent to a complex multiplication on a digital machine, which implies a raw computation rate roughly equivalent to 10^{12} FLOPS (although this does not account for efficient algorithms that can be used on the

digital machine). For the longest replica, the correlator computes at the peak rate of 563×10^9 multiplications per second but yields only one correlation point, an inefficient mode of operation. Typical values of T_{REF} in our system are between one-third and two-thirds of the optical aperture.

The last two columns of Table 2 compare the maximum average number of correlation function samples produced per second from Eq. (10) with the number of FLOPS required from an equivalent digital processor to equal this rate from Eq. (19). These equations are also plotted in Fig. 4. We have used the values $N_{REF} = 1024$ and $N_{opt} = 8192$ in Eq. (19); the results vary by roughly ten percent for values of N_{REF} smaller or larger by a factor of two, so that these values are reasonable for the range of T_{REF} used in our system. To produce on average approximately 50 million correlation samples per second, the equivalent digital processor would require approximately 2 GFLOPS of power, although this should be increased to 8-10 GFLOPS to account for typical inefficiencies in a digital machine.




Fig. 4.—Comparison of the maximum average number of correlation function samples produced per second, Eq. (10), with the computation rate required of an equivalent digital processor, Eq. (19), as a function of τ_{REF} . For this figure, $f_{CLK} = 80$ MHz, $N_{REF} = 1024$, and $N_{opt} = 8192$.

Input and Output Bandwidth Requirements

To make use of the full potential of the AO correlator, the I/O interface between the host computer and the optical correlator must supply the raw material (input waveforms) and accept the product (correlation functions) at the rates required by the optical correlator. Without this condition, T_{OVR} can never be small. The maximum I/O bandwidths essentially depend on the bandwidth of the AO cells, although the bandwidth manifests itself in the equations below as a clock rate. We begin by deriving the maximum required I/O bandwidths, and then we discuss the implications of these rates.

The greatest bandwidth requirement for input data occurs when a new reference waveform and a new return waveform are used for every correlation. The input bandwidth requirement is given by the ratio of the number of waveform samples to the time required for one correlation sweep,

$$f_{IN} = \frac{f_{CLK}(T_{RCV} + T_{REF})}{T_{RCV} + T_{OVR}} = \frac{2f_{CLK}}{2 - \tau_{REF} + \tau_{OVR}} \quad (20)$$

in samples per second. The maximum input bandwidth f_{IN}^{max} occurs when the overhead is zero.

The largest output bandwidth requirement occurs when every correlation sweep is digitized and returned to the host. This results in an output bandwidth requirement of f_{OUT} samples per second given by

$$f_{OUT} = \frac{(2f_{CLK})(\frac{1}{2}T_{COR})}{T_{RPT}} = \frac{2f_{CLK}(1 - \tau_{REF})}{2 - \tau_{REF} + \tau_{OVR}}, \quad (21)$$

where we have assumed that the output sample rate is twice the clock rate of the input waveform generator D/A converters. Again, the maximum output bandwidth f_{OUT}^{max} occurs when the overhead is zero.

Figure 5 displays the bandwidth requirements from Eqs. (20) and (21) for our system for $\tau_{OVR} = 0$. Note that the input bandwidth is large for all values of τ_{REF} , but the output bandwidth decreases with τ_{REF} . This is an indication of the diminishing efficiency with increased replica length. If the input samples are in a typical floating point format (eight bytes per complex sample), the input bandwidth ranges from 640 to 1280 MB/s. If one were to convert the samples to two-byte complex integer samples, the required data rate would be cut by a factor four, although converting floating point numbers to integers at this rate requires considerable computational resources. Note that the sum of the maximum input and output bandwidth requirements is a constant,

$$f_{IN}^{max} + f_{OUT}^{max} = 2f_{CLK}. \quad (22)$$

With typical values for f_{CLK} and common data formats, these data rates will overwhelm all but the fastest of today's I/O systems, although particular applications may have less severe requirements, as shown below.

DESIGN GOALS

The primary goal of the optical correlator project was to demonstrate as much of the capability of the optical system as possible by installing and testing it in existing digital signal processing systems. The systems used as testbeds were built around medium- to high-end commercial hardware, but even the latter imposed restrictive I/O bandwidths. Because of budget constraints and the limited performance of available electronic technology, the electronic interface and control system for the optical correlator were designed to realize only a fraction of the capability of the optics. Nevertheless, building even a limited system allowed us to demonstrate the concept, test the performance of the optics, obtain a measure of the performance enhancements attainable with AO technology, and determine the necessary improvements for the electronic interface. With this goal in mind, we chose an application that could demonstrate the capability of the acousto-optical system without imposing severe I/O bandwidth requirements. We also required that the system interface with a variety of host systems, and that the system be easily accessed from a variety of high-level languages through calls to a subroutine library. We addressed the requirements in the following ways:




Fig. 5.—Input and output data rates from Eqs. (20) and (21) in samples per second as a function of replica duration for the case $f_{CLK} = 80$ MHz and $\tau_{OVR} = 0$.

1. The application we chose required that the return waveform be cross-correlated with multiple time-scaled (e.g., Doppler-compensated) variants of a replica of the transmitted wide band waveform. We derived the variant replicas from the original replica by scaling the time variable, which was effected by reading out the original replica with an adjustable clock. (This method of waveform generation also partially satisfied the wide band Doppler-compensation¹¹ requirement.) Thus, with only a single replica (or a small "library" of distinct replicas) and a return waveform sent to the optical correlator system, many (up to a few hundred) correlation functions could be generated. This effectively reduced the input bandwidth by the required number of Doppler variants while still taking advantage of the optical correlator's capabilities.
2. In Doppler processing of wide band signals, a simple shift of the carrier is not enough to compensate for the Doppler effect; rather, time warping or scaling must be used. Because the input waveforms to the optical correlator in our application have been demodulated from a carrier, an additional shift of the AO cell carrier is usually required to obtain maximum correlation. We satisfied this by providing an adjustable carrier for one of the quadrature modulators, as described below.
3. The output bandwidth was reduced by including a constant-false-alarm rate (CFAR)¹² circuit to detect threshold crossings. The CFAR circuit compares the correlation function with a threshold function: whenever the correlation function exceeds the threshold, a report is generated and forwarded for further processing. Depending upon the false-alarm rate, the output bandwidth can typically be reduced by a factor of 10 to 100. In an alternate mode, the entire correlation function, in a sampled and digitized format, can be returned to the host, but at the cost of reduced throughput since the bandwidth supported by the interface hardware in our system is much less than that required by Eq. (21).

4. We designed the system to interface with a variety of computers, including VAX computers and with IBM-compatible PC's. The PC was used as a development, testing, and demonstration platform. The heart of the correlator system interface is a 16-bit parallel first-in first-out (FIFO) buffer that is easily interfaced to a wide variety of computers.
5. The logical interface consists of two sets of software: one set resides on the host, and consists of a library of subroutines written in FORTRAN; the other set runs on the optical correlator System Controller and responds to the commands and data issued by the host software. The host software (on a VAX) is callable from C, FORTRAN, Pascal, or Ada.
6. Typically, the return waveforms in an application are too long to be processed in one correlation sweep, so the system provides for overlap processing of the return waveforms. The analogous restriction arises in a digital system due to the finite size of the FFT processor. The optical system can actually process a long waveform continuously by reading out the return waveform without interruption and correlating it against suitably delayed copies of the replica. The application we chose does not lend itself to this technique, however.
7. The system accepts input data in a variety of formats, including 32-bit floating point, one's- or two's-complement integers, and unsigned integers. This requirement allows the processor to coexist with most digital processors. The system also returns data in a format acceptable to the host.
8. Data formats are directly related to dynamic range capabilities. The correlator system has three different dynamic ranges: input, internal, and output. The internal dynamic range limits are set by the optical and electronic hardware used in the system. In particular, the Waveform Generator D/A converters require eight-bit integers. The requirement for flexible data formats necessitates normalization of the data followed by conversion into an integer format for the D/A converter. The linear output dynamic range is limited by the internal dynamic range and by the resolution of the A/D converter.

SYSTEM ARCHITECTURE

The block diagram of Fig. 6 shows a high level view of the signal processing system. An array of sensors (or an antenna) collects the signal and noise that are present in the signal channel. The analog output of the sensors is then filtered, sampled, and digitized. These samples are then processed to form return waveforms. For example, this process might consist of spectral or spatial filtering (beamforming) to enhance the signal-to-noise ratio. The return waveforms are then transferred to the wide band optical correlator, where they are crosscorrelated with a replica of the desired waveform. The correlator output is then transferred to the host system where it is put into a form suitable for display and interpretation. The hardware and software for the system are described in more detail in the following sections.

Hardware

The hardware contains two major assemblies: an optical system and an electronic system. The optical system contains the basic computational engine; it operates on analog signals and produces an analog correlation function. The electronic system provides the signals to drive the acousto-optical cells, accepts the output of the optical detector, provides the interface between the optical system and the digital host system, and executes command requests from the host. Because of the high clock rates and large bandwidths needed for the system, the electronic system required a significant development effort that was

hampered at the time by limitations in the resolution of the A/D and D/A converters at the speeds required, and by the unavailability of memory of sufficient density and speed.

Optical System

The optical engine consists of a one-dimensional, space-integrating, optical heterodyne detection correlator.^{1,13} The component layout is shown in top and side plan views in Fig. 7. An Hitachi 8312E laser diode laser source (20 mW at 830 nm) and beam-shaping optics produce a collimated beam of light with a cross-section approximately 31 mm by 2 mm. A portion of this beam is diffracted by the acoustic waveform in the first AO cell (operating in the Bragg regime), thus imparting the phase and amplitude information of the waveform to the optical beam. The sliding action or lag required for convolution or correlation is provided naturally by the propagation of the acoustic signal along the length of the Bragg cell. The diffracted beam then propagates a short distance to the next cell where it is diffracted a second time. This doubly-diffracted light thus contains the product of the two waveforms. The integration of this product of waveforms is performed by a lens that focuses the doubly-diffracted beam and the remaining undiffracted beam onto a photodiode. The photodiode acts as a heterodyne detector that converts the high frequency optical information into an electronic signal for further processing.




Fig. 6.—High level block diagram of the signal processing system.

The Bragg cells are shear-mode TeO₂, Model No. N45075-6-20, manufactured by Newport Electro-Optic Systems with a length of 75 μ s, acoustic direction [110], optical direction [001], and a diffraction efficiency of 50%/W. The cells were specially made to minimize the separation between the acoustic beams: the center-to-center separation in our layout is approximately 5 mm. At lower time-bandwidth products (less than 200), this "shadow casting" method works well, but distortion in the correlation signal appears at higher time-bandwidth products. The values for the center frequency f_{cell} , the bandwidth W_{cell} , and the optical aperture (or useful length) T_{OA} of our cells are shown in Table 3.

The Bragg cells are shear-mode TeO₂, Model No. N45075-6-20, manufactured by Newport Electro-Optic Systems with a length of 75 μ s, acoustic direction [110], optical direction [001], and a diffraction efficiency of 50%/W. The cells were specially made to minimize the separation between the acoustic beams: the center-to-center separation in our layout is approximately 5 mm. At lower time-bandwidth

products (less than 200), this "shadow casting" method works well, but distortion in the correlation signal appears at higher time-bandwidth products. The values for the center frequency f_{cell} , the bandwidth W_{cell} , and the optical aperture (or useful length) T_{OA} of our cells are shown in Table 3.




Fig. 7.—Top and side views of the optical layout.

The focusing or integrating (Fourier transform) lens is a laser diode glass doublet Model 06LAI013/076, from Melles Griot. Its focal length is 145 mm at 830 nm and approximately 75 percent of its 40 mm aperture is used. The lens is anti-reflection coated for a maximum reflectance of 0.25% in the region 780-830 nm. The wavefront distortion is better than $\lambda/5$, but with a small defocus the design appears to be capable of four times better performance ($\lambda/20$). In practice, the photodiode detector is experimentally positioned for optimum response to a test signal. The algorithm used to optimize the detector position attempts to maximize the useful optical aperture, uniformity of response across the field, and signal-to-noise ratio.

The optical signal is detected by a Hewlett-Packard Model 5082-4205 photodiode with an effective area of $3 \times 10^{-3} \text{ cm}^2$. The photodiode is mounted directly onto an Analog Modules 713A preamplifier. The preamplifier output is connected to an RHG Electronics Laboratory, Inc. ICLT150B log amplifier. The log amplifier has a center frequency of 150 MHz, a bandwidth of 100 MHz, and a video risetime of less than 10 ns. The dynamic range extends from -70 dBm to +5 dBm.

Table 3—Acousto-Optical Cell Parameters

Parameter	Value
f_{cell}	75 MHz
W_{cell}	40 MHz
T_{OA}	44 μ s

Electronic System

The electronic hardware for the wide band correlator is housed in a standard nineteen-inch crate designed to hold printed circuit boards or cards meeting the standard 6U Eurocard form factor. The cards are connected via 64-pin connectors to a custom backplane. Power, control, and data lines are all provided via the backplane. A block diagram of the system is shown in Fig. 8. The crate contains seven boards or sub-assemblies and a custom power supply. The sub-assemblies include a System Controller board, two Waveform Generator boards, two Clock and Carrier boards, a Quadrature Modulator and Amplifier board, and a Constant-False-Alarm Rate board.

The System Controller board consists of the Control Computer, dynamic random access and programmable read-only memory, interface logic for the backplane, and a bi-directional FIFO buffer and associated circuitry for connection to the host. The Controller is responsible for most of the timing and logic of the system, and is responsible for all communications with the host. The Control Computer is based on a Texas Instruments TMS320C25 digital signal processing microprocessor with 64 KB of memory for data and 64 KB for code. The system clock frequency is 36 MHz, which makes the microprocessor operate at 9 MIPS. The controller is programmed in C and assembly language; the code resides in UV-erasable PROM's. The software executing on the Control Computer acts upon command requests issued by the host, and returns the status and results of the command execution. The Controller can also perform various diagnostic and setup functions.

Data transfer to and from the host system is via a first-in first-out dual channel buffer that is 17 bits wide by 16 words deep. The system can be interfaced to several popular buses, including the PC bus, the VAX UNIBUS, and the VAXBI bus. Data transfer rates with the PC are limited to a few tens of KB/s. A DEC DR11-W card is required for communication via the VAX UNIBUS. This card is limited to sustained data transfer rates of a few hundred KB/s. The VAXBI bus is capable of transferring between one-half and one MB/s using the DEC DRB32-W parallel interface card.

Two Waveform Generators, each consisting of a digital waveform buffer (Replica or Signal), two eight-bit digital-to-analog converters, a quadrature modulator, and an RF amplifier, produce the analog replica and return waveform signals that drive the Bragg cell transducers. The target applications for the correlator system use quadrature (or Hilbert transform) demodulation to reduce the required sample rate; thus, the return and replica waveforms consist of complex samples. The real and imaginary components are connected to the I and Q inputs of IntraAction MD-75W quadrature modulators that modulate a nominally 75 MHz carrier. The A/D converters operate at a nominal conversion rate of 80 MHz. The 16 KB return waveform (Signal) buffer more than accommodates the longest waveform supported by the Bragg cells. The 256 KB replica waveform buffer can store multiple replicas: the particular waveform used for a crosscorrelation is selected via software at correlation time. To minimize power consumption,

the memory is divided into four banks of CMOS memory operating at 20 MHz. These banks are multiplexed to give an effective clock rate of 80 MHz.




Fig. 8.—Block diagram of the optical correlator system.

Two clock signals and two carrier waves are needed for the system. A Vectron CO-233ME 80 MHz ECL clock drives one buffer, and a Vectron CO-233 75 MHz sine wave oscillator provides the carrier for the associated quadrature modulator. The second buffer memory requires an adjustable clock and carrier to perform the time-scaling operation. An 80 ± 6 MHz ECL clock and a 75 ± 6 MHz sine wave are produced by a direct digital synthesis (DDS) technique. Spurious signals in the DDS system are down at least 55 dBc across the band. Switching time for a frequency accuracy of one part per million is less than 2 μ s.

The Threshold Comparison or CFAR unit consists of an eight-bit A/D converter, an eight-bit D/A converter, two blocks of memory called COR and THR, and a comparator circuit. The CFAR circuitry runs at one-half the clock rate of the waveform generators, producing a decimation by one-fourth (because of the counter propagating waveforms) of the correlation function samples. In the digitize mode, the CFAR unit samples and digitizes the correlation functions, and stores the samples in COR memory for uploading to the host or for use in the comparison mode.

In the detection or comparison mode, the CFAR unit uses the D/A converter to generate an analog threshold function from samples in THR memory. The threshold function samples are either downloaded by the user or computed internally by the System Controller from one of several user-selectable algorithms. The algorithms are essentially two-pass normalizers¹⁴ that attempt to estimate the noise in a waveform that is assumed to contain noise and signal. These internally computed threshold functions can also be uploaded to the host at the user's request. During a correlation sweep in the comparison mode, the threshold function is generated synchronously with the correlation function and provides one input to the comparator circuit; the other input is the analog correlation function. At each clock pulse, whenever the level of the correlation function exceeds the threshold function level, the comparator signals that the correlation function should be sampled, digitized, and stored in COR memory; otherwise a zero is written to COR memory. The threshold crossings can then be uploaded to the host as a collection of threshold crossing messages. A threshold crossing message consists of the correlation amplitude, the threshold amplitude, and the time (sample number) of the crossing. If the threshold function was generated by a CFAR algorithm, the output data bandwidth (threshold-crossing messages per second) is statistically predictable. For typical threshold settings, this method can reduce the output data bandwidth by one or two orders of magnitude compared to returning the digitized correlation functions.

Software

The correlator is designed to operate as a slave processor to a master host system. From the user's point of view, the opto-electronic correlator appears much like a vector or array processor that is accessed via a set of high level language subroutine calls. To access the correlator, the user writes application code that calls the desired subroutines, links his compiled application code with the correlator subroutine library, and then runs the application. Once the correlator hardware and software are installed on the host system, the opto-electronic nature of the correlator hardware is essentially transparent to the user.

The correlator subroutines¹⁵ are divided into three classes: *Setup*, *Signal Processing*, and *Diagnostic*. The *Setup* subroutines perform initialization tasks such as device allocation and calibration. The *Signal Processing* subroutines set sampling rates, download waveforms to the correlator, perform correlations and upload them, generate threshold functions, and perform detection processing. The *Diagnostic* subroutines allow the user to test the optics and electronics, as well as access the functionality of the correlator at a lower level for specialized processing or diagnostic purposes.

Each subroutine consists of two modules: one residing on the host system, the other residing on the correlator. Each host module checks the user-supplied data for consistency, formats it, and then transfers the data along with a control code to the System Controller. The correlator module then completes the consistency checks, executes the action requested by the control code, and then formats and returns the results to the host. The System Controller maintains a parameter database that stores essential information about the replica and return waveforms, overlap information, user specified parameters, and the state of the system, including internal timing and calibration parameters. To maintain the integrity of the database, a user may call the *Setup* and *Signal Processing* subroutines only in a specified order. This is enforced with a "level" mechanism that assigns a level number to each of these subroutines.

There are two *Signal Processing* subroutines that generate the correlation functions. The first, *OBC_BKGD_CORR*, corresponds to the digitize mode of operation of the CFAR circuit. It crosscorrelates the replica and return waveforms for up to eight values of the Doppler estimate, digitizes and stores the results, and (optionally) returns copies of the correlation functions to the host. The second corresponds to the detect or comparison mode of the CFAR circuit. It correlates the return waveform against all Doppler-compensated replicas across a user-specified range from a low Doppler bin to a high

Doppler bin. It compares each correlation function with a threshold function; any threshold crossings are stored temporarily and then transferred in batches to the host.

Since the optical correlator requires eight-bit integer waveform samples to drive the D/A converters, a major function of the software is to normalize the input waveforms to match their dynamic range as well as possible with that of the DAC's. This is done automatically as block normalization by the subroutines that transfer the replica and return waveforms to the correlator. In the case of a replica waveform, the block consists of the entire waveform; for that of a return waveform, a block is the current section of the waveform sent by the user to the correlator. The normalization method used for each type of waveform is selected by the user with a Setup subroutine. Some of the possible methods are limiting, linear scaling, and logarithmic scaling. Each of these methods can normalize by a user-supplied value, by the peak amplitude in a block, or by other methods based on the "local" statistics of the data. A modified mu-law method¹⁶ that we have developed compresses the amplitude logarithmically but preserves the phase. All of the methods preserve the phase as much as possible, since, for our waveforms, it is the phase information that provides most of the signal processing gain.

When connected to a VAX computer, the user has the option of linking the same application code to a library of emulation subroutines rather than to the optical correlator subroutine library. These subroutines emulate the functions of the optical correlator. Two versions of the emulator library are available: a version that uses the standard VAX scalar processor, and a version that uses the VAX vector processor (if it is available on the user's VAX). Since these subroutines execute directly on the VAX, they have no I/O overhead associated with transferring data over the parallel interface to and from the correlator (although there may be I/O overhead associated with virtual memory). Since the emulator subroutines are functionally identical to the optical correlator subroutines, they provide a convenient means to compare the performance of the digital (emulator) and optical correlators.

RESULTS

We conducted a series of performance tests of the optical correlator connected either to a PC or to a VAX. The PC was a Compaq SLT 386S/20 portable microcomputer attached to an expansion unit. Communication with the correlator was via a 16-bit I/O port over the PC bus. The VAX was a VAXVector 6410, that is, a VAX 6410 with an FV64A vector processor installed. Data communication between the VAX and the correlator was via a DEC DRB32W 16-bit parallel interface card. Neither of these machines can provide the output bandwidth required by Eq. (21). As shown below, the electronic interface and control system takes advantage of less than one percent of the available power of the optical correlator engine. Even so, the correlator was typically twenty to seventy times faster than an FFT correlation algorithm using the VAX vector library routines. Thus, an efficient implementation of an optical correlator system could perform several thousand times faster than the VAXVector 6410.

I/O Bandwidth Performance

As shown in Fig. 5, the I/O bandwidth required to keep the optical system fully occupied is potentially quite large. We lowered the input data rate in our system by taking advantage of two features of the application: the replica waveforms need be updated only occasionally, and the return waveform section must be updated only after it is crosscorrelated with many time-scaled replicas. The time-scaled replicas are generated on demand, in real time, from the unscaled replica by the Waveform Generator driven by the variable clock. Since the replica waveforms change only occasionally, they may be downloaded once and stored in the correlator system. If we ignore the small contribution of the replica waveform to the input bandwidth requirement, and if we denote the number of time-scaled replicas by N_{Dpl} , Eq. (20) becomes

$$f_{IN} = \frac{f_{CLK}(2 - \tau_{REF})}{N_{Dplr}(2 - \tau_{REF} + \tau_{OIR})} \quad (23)$$

For example, if $N_{Dplr} = 100$, the maximum input bandwidth requirement varies from 0.8 to 1.6 MSPS (mega-samples per second), a requirement that is much easier to satisfy.

Lowering the output data rate is only possible by additional processing of the sampled correlation functions. Otherwise, the full output bandwidth must be supported to make maximum use of the correlator. Since in our application the correlation functions are ultimately used to generate detections, only the detections need be reported to the host if they can be generated on the correlator system. Our system does this by comparing each correlation function with a threshold function in the CFAR circuit, as described previously. Setting the threshold level to produce a false-alarm rate of one percent can lower the output data rate by approximately two orders of magnitude.

Our system transferred data at a rate that only approached 1 MB/s. The data rates were limited by the I/O hardware on the host computer and by an inefficient memory management scheme in the correlator system. The implications of this limitation are described in the next section.

Computational Performance Measurements

The efficiency parameter of Eq. (12) provides a measure of the effect of overhead on the performance of the optical processor. The overhead time depends on the hardware setup time, on the time required to transfer data to or from the host during a correlation period, and on the amount of code that must be executed by the System Controller between each correlation period. The correlator supports three subroutines that perform correlations. The first is a diagnostic routine that continuously correlates a test waveform against a zero-padded copy of the waveform, i.e., it autocorrelates the test waveform. The other two subroutines have been described previously in the Software section; we refer to them as the digitize and the detect subroutines. We used these subroutines to measure the efficiency of the correlator.

The diagnostic correlation subroutine contains a small amount of code that performs a simple computation, sets up the various registers and counters, initiates the correlation, and then waits for the end-of-correlation signal before beginning again. It correlates a 2048 sample waveform against a simulated return waveform consisting of 2048 zeroes, concatenated with a copy of the waveform, concatenated with enough zeroes to make the return waveform satisfy Eq. (5). Thus, the return waveform effectively contains 4992 samples. Under these conditions, the observed correlation period was $T_{RPT} = 189 \mu s$, which implies an overhead time of $127 \mu s$. Of this, about $77 \mu s$ was attributable to code execution, about $35 \mu s$ was due to padding the return waveform with an excessive amount of zeroes, and about $15 \mu s$ was due to an inefficient algorithm for filling the AO cell "pipeline," i.e., the software waited until a correlation was completed before it instructed the hardware to begin converting waveform samples for the next correlation cycle. The resulting values for Eqs. (9), (10), and (12) are shown in Table 4. The efficiency of 33 percent indicates the necessity for asynchronous software, i.e., software that executes during the correlation. Even with asynchronous operation, the software efficiency must be improved, and the code probably must be executed by a faster microprocessor. Furthermore, no I/O operations with the host occur during this subroutine.

Table 4—Performance of the Continuously Correlating Diagnostic Subroutine

Conditions:	$T_{REF} = 25.6 \mu\text{sec}$ (2048 samples) $T_{RCV} = 62.4 \mu\text{sec}$ (4992 samples) $T_{RPT} = 189 \mu\text{sec}$	
Avg. no. correlation samples/s, Eq. (9)	f_{avg}	15.6 MSPS
Max. avg. no. correlation samples/s, Eq. (10)	f_{avg}^{max}	47.2 MSPS
Efficiency, Eq. (12)	η	0.33

Measurements made on the OBC_DETECT and OBC_BKGD_CORR correlation subroutines indicate that they are operating at less than one percent efficiency. This efficiency refers to the time between consecutive correlations; it does not include the time required to transfer data to the host. If it did, the efficiency would drop further by factors of approximately 2 to 10, depending upon whether the host was a VAX 6410 or IBM PC.

Table 5—Comparison of optical, vector, and scalar correlators. Errors are one standard deviation.

No. Correls.	t_{Opt} (ms)	t_{Vector} (ms)	t_{Vector} / t_{Opt}	t_{Scalar} (ms)	t_{Scalar} / t_{Opt}
1	18±9	320±23	18	991±18	55
2	19±8	714±40	38	2021±39	106
4	26±9	1415±43	54	4232±119	163
6	32±10	2136±56	67	6274±185	196
8	41±12	2767±60	67	8294±198	202

In another set of tests, we compared the performance of the optical correlator connected to the VAX against the two emulation correlators: an FFT correlator using the VAX vector processor, and an FFT correlator using the normal VAX scalar processor. The same data sets were run multiple times on all three correlators using the digitize subroutine. The results are shown in Table 5 and Fig. 9. The table lists the average CPU time charged to the subroutine for each of the three correlators as a function of the number of correlation sweeps performed per subroutine call. It also compares the times for the two emulator subroutines with the optical correlator. The uncertainties in Table 5 indicate one standard deviation: they are due primarily to activity from other processes on the VAX (running VMS). We also monitored the number of page faults (disk accesses necessitated by too little memory) for the process. Typically, the number of page faults was significant for the first run, but was zero or one for subsequent runs. (The total number of runs was typically 75 to 100.) For the digital correlators, the timing uncertainties are relatively insignificant, but for the optical correlator they are large since the time used by the correlator is small.

Figure 9 plots the data in Table 5 on a semi-log scale to emphasize the differences in performance. One should note that the figures for the optical correlator are more than two-hundred times greater than the theoretical minimum time given by Eq. (7), indicating an efficiency of less than half a percent.




Fig. 9.—Comparison of scalar and vector digital processors with the optical correlator as a function of the number of correlation sweeps. Times for the optical correlator include the time required to generate and digitize the correlation function and to upload it to the host computer.

Dynamic Range Measurements

The linear dynamic range of acousto-optical cells¹⁷ is limited primarily by intermodulation distortion. The distortion can occur in the electronics driving the AO cells, from multiple diffractions within a cell, or from nonlinear piezoelectric and acoustic effects. If one is more interested in detections than concerned about distortions in the correlation function, the useful dynamic range of an AO correlator system may extend well beyond the linear dynamic range of the AO cells.

To measure the dynamic range performance of the correlator, we performed two tests. In the first test, a sinusoidal wave of duration $\tau_{REF} = 0.85$ was autocorrelated. Prior to downloading, the amplitude of the digital reference waveform was attenuated numerically. The return waveform consisted of a section of leading zeroes, an unattenuated copy of the reference waveform, and a section of trailing zeroes. The peak voltage of the correlation function was measured on an oscilloscope at the (buffered) output of the log amplifier. The relative drop in peak height for various levels of attenuation is shown in Table 6. The response is reasonably linear only for the first two decades of attenuation. From our experience with the analog quadrature modulators and RF power amplifiers, we postulate that the non-linearity is due primarily to the quadrature modulators: as the input signal level decreases, the non-linearity of the balanced mixer increases. A portion of the non-linearity is probably attributable to the acousto-optical system, as shown by the following test.

Table 6—Decrease in Peak Height as a Function of Waveform Amplitude. The reference waveform was attenuated digitally prior to downloading to the optical correlator system. The return waveform was not attenuated.

Attenuation (dB)	Relative Peak Height (dB)
6.02	-5.32
20.0	-19.3
26.0	-24.4
32.0	-29.1
40.0	-33.0
42.1	-33.3




Fig. 10.—Dynamic range performance of the optical correlator system.

In the second test, a full-scale waveform identical to that used in the first test was autocorrelated, but the output of each RF power amplifier was attenuated prior to the AO cell transducers. The attenuation was applied with Hewlett Packard Model 355D RF attenuators to the reference waveform, to the return waveform, or to both waveforms as shown in Table 7. Linearity of the quadrature modulators and amplifiers was not an issue, since their average power levels remained constant. For this test, the response is approximately linear for at least three decades of attenuation. The improved results for this test indicate that the quadrature modulators or power amplifiers are responsible for the poorer performance in the first test.

The combined results of the first and second tests are plotted in Fig. 10, where the results of the first test are indicated by open squares, and the results of the second test are indicated by triangles. We have not done an exhaustive search to determine the best combination of drive levels for the quadrature modulators and power amplifiers, although our experience to this point leads us to expect only modest improvements by this approach. A more effective approach would be to replace the existing quadrature modulators with improved ones, or to use a digital quadrature modulation technique.

Table 7—Decrease in Peak Height as a Function of AO Cell Power. First column is the sum of the attenuation applied to each cell. The second column contains results for the attenuation applied only to the reference waveform, and the third column for the return waveform only. The fourth column contains results for attenuation applied to both waveforms.

Total Attenuation (dB)	Relative Peak Height		
	Reference (dB)	Return (dB)	Reference and Return (dB)
10.00	-9.2	-9.8	-
20.0	-19.6	-19.9	-19.3
30.0	-29.1	-29.1	-28.7
40.0	-36.7	-36.1	-36.4
50.0	-42.0	-	-

CONCLUSION

We have described an acousto-optical correlator system that we developed, we have derived performance measures for it that allow comparison with conventional digital correlators, and we have described the results of several performance tests. We described our design goals for the system, and discussed how the goals influenced the system. In terms of computation speed, the optical correlator is potentially hundreds to thousands of times faster than currently available digital machines. When volume and power considerations are important, the optical system is without peer. Even though its current electronic interface package limits its performance to less than one percent of its potential, it performed 20 to 70 times faster than a VAX 6410 using a vector processor and an optimized FFT correlation routine. A two-dimensional optical system has the potential to enhance the computation rate by a factor of 25 to 100. The challenge to integrating such a powerful computational engine into a digital signal processing system is due primarily to the large input and output data bandwidths, but also to the necessary data transformations between the analog and digital domains.

The current system has several shortcomings: a) The internal control hardware and software are inefficient: their performance needs to be improved by nearly a factor of two hundred. b) Even with efficient control the input and output bandwidths are limited to less than 1 MB/s. A high performance I/O interface such as HiPPI or FDDI would make the correlator available for a much wider set of applications.

c) The eight-bit resolution of the Waveform Generator D/A converters limits the correlator to applications where normalization of the input waveforms is acceptable. This is not a severe restriction, but it does impose an additional computational burden because normalization of the waveforms is required. d) The performance of the analog quadrature modulators is not adequate even for the eight bit waveforms. A digital quadrature modulation scheme can remove this restriction. e) The eight-bit resolution of the output A/D converter is a more serious limitation. It will only be removed when faster high-resolution converters become available. f) The system only processes the amplitude of the correlation function. The ability to report the phase information as well requires a more complex detection system. g) The shadow casting method used to transfer the optical signal from the first AO cell to the second is inadequate at higher time-bandwidth products. This restriction can be removed with a high performance relay lens system. We are addressing all of these limitations in a second generation optical correlator system.

ACKNOWLEDGMENTS

The authors wish to thank Scott Adams for his assistance with designing, implementing, and conducting some of the performance studies. This work was sponsored by the Office of Naval Research. Much of the hardware and software for this system were built by the Westinghouse Science and Technology Center under contracts N00039-86-C-0060 and N00014-87-C-2058.

REFERENCES

1. A. VanderLugt, *Optical Signal Processing* (John Wiley & Sons, Inc., New York, 1992).
2. N. J. Berg and J. N. Lee, eds., *Acousto-Optic Signal Processing*, (Marcel Dekker, New York, 1983).
3. *Proceedings of the SPIE 1993 Conference on Transition of Optical Processors into Systems*, D. P. Casasent, ed. , Vol. 1958, (SPIE, Bellingham, WA, 1993).
4. A. D. Whalen, *Detection of Signals in Noise* (Academic Press, Orlando, 1971), Ch. 6, p. 167.
5. J. N. Lee, ed., *Design Issues in Optical Processing*, (Cambridge Univ. Press), Chap. 2, in preparation.
6. *Opt. Eng.* 31 1821-1916 (1992), special section on wavelet transforms.
7. R. E. Crochiere and L. R. Rabiner, *Multirate Digital Signal Processing* (Prentice-Hall, Englewood Cliffs, NJ, 1983).
8. E. O. Brigham, *The Fast Fourier Transform* (Prentice Hall, Englewood Cliffs, NJ, 1974).
9. O. D. Grace and S. P. Pitt, "Sampling and interpolation of bandlimited signals by quadrature methods," *J. Acoust. Soc. Amer.* 48, 1311-1318 (1969).
10. H. D. Helms, "Fast Fourier transform method of computing difference equations and simulating filters," *IEEE Trans. Audio and Electroacoustics* AU-15, 85-90 (1970).
11. E. J. Kelly and R. P. Wishner, "Matched-filter theory for high-velocity, accelerating targets," *IEEE Trans. Military Electron.* MIL-9, 56-69 (1965).
12. M. I. Skolnik, *Introduction to Radar Systems* (McGraw-Hill, New York, 1962), Ch. 10., pp. 392-395.

13. H. H. Arsenault, T. Szoplik, and B. Macukow, eds., *Optical Processing and Computing*, (Academic Press, San Diego, 1989).
14. W. A. Struzinski and E. D. Lowe, "The effect of improper normalization on the performance of an automated energy detector," *J. Acoust. Soc. Am.* **78**, 936-941 (1985). D. V. Gupta, J. F. Vetelino, T. J. Curry, and J. T. Francis, "An adaptive threshold system for nonstationary noise backgrounds," *IEEE Trans. Aerosp. Electron. Syst. AES-13*, 11-16 (1977).
15. R. D. Griffin, A. R. Schuler, J. N. Lee, "NRL Optical Broadband Correlator Subroutine Reference Manual: Version 2.0," NRL Memorandum Report NRL/MR/6523—92-6972, April 1992.
16. B. Smith, "Instantaneous companding of quantized signals," *Bell Syst. Tech. J.* **36**, 653-709 (1957).
17. G. Elston, "Intermodulation products in acousto-optic signal processing systems," *Proceedings of 1985 Ultrasonics Symposium*, (1985), pp. 391-397. D. L. Hecht, "Multifrequency acoustooptic diffraction," *IEEE Trans. Sonics Ultrasonics SU-24*, 7-18 (1977).